arXiv:1603.01607v1 [cs.DS] 4 Mar 2016

COMPUTING SHORTEST PATHS USING A*, LANDMARKS, AND POLYGON INEQUALITIES

Newton H. Campbell Jr., Raytheon BBN Technologies

SIAM Workshop on Network Science 2015
May 16-17 · Snowbird, UT

Summary

We introduce a new dual-landmark heuristic for the A* algorithm that references a data structure of size $\theta(|\boldsymbol{L}|^2 + |V|)$, where $\boldsymbol{L}$ represents a set of strategically chosen landmark vertices and V the set of vertices in the graph. This heuristic's benefits are permitted by a new approach for computing lower bounds based on generalized polygon inequalities, in which each landmark stores the distances between it and vertices within its graph partition. In this paper, we demonstrate experimental results that show exemplify the space complexity reduction and speedup in comparison to its single-landmark counterpart for large graphs.

Introduction

Modern navigation planning requires the ability to regularly compute the shortest path between two points in massive road networks. In such cases, preprocessing algorithms are used to increase the performance of shortest path queries. Many such algorithms require heavy upfront computation and storage. Few algorithms concern themselves with the space complexity required to aid queries. The problem that this research addresses is that modern shortest path preprocessing algorithms have space and preprocessing time requirements for large-scale graphs that are impractical for real-world applications.

A*, Landmarks, and Polygon Inequalities

ALT [2] describes a preprocessing technique for shortest path queries that, prior to query time, chooses a relatively small number of landmark nodes in a graph and computes the distances between all vertices and these landmarks, allowing the A* algorithm to leverage the triangle inequality during search queries. The algorithm works as follows: For a simple graph G with vertices $A, l, C \in V$, where l is a landmark vertex chosen beforehand, the shortest path distances between each vertex allow the graph to form a metric space. Therefore, for the distances between vertices $A, l, C \in V$, the following reverse triangle inequality holds:

$$|d(A, l) - d(l, C)| \leq d(A, C) \tag{1}$$

ALT uses this inequality to create a heuristic for A* that estimates the distance between vertices A and C. By computing and storing the values between each chosen landmark and all vertices in the graph prior to performing any shortest path queries, this estimate is performed for each computed landmark upon a visit to vertex A. The maximum of these estimates is the heuristic function's value, denoted as π_t^L.

By using information about multiple landmarks, new lower bounds can be computed from either generalized polygon inequalities or ones specific to any shape embedded within the graph. The use of these new lower bounds as a heuristic has resulted in a new class of algorithms called ALP, for A*, Landmarks, and Polygon Inequalities. The base case for this class of algorithms is the heuristic used for the ALT algorithm. Here, we demonstrate that polygon inequalities for quadrilaterals can also be used to establish the lower bounds for shortest path queries in a graph. The reverse quadrilateral inequalities shown in Figure 1 hold for a graph G with source and target vertices $s, t \in V$ and chosen landmarks $l_1, l_2 \in V$. The first five are derived from the triangle inequality as applied to quadrilaterals. A potential problem with these inequalities is that they have ability to generate negative lower bound estimates. However, because multiple points are used, a varying set of inequalities can be generated to estimate distances. When attempting to estimate lower bounds using ALP, other inequalities should be considered such that the highest possible estimate can be used. We use the sixth equation, derived from Ptolemy's inequality, as a concrete example for quadrilaterals.

Just as with ALT, the maximum over the set of these lower bounds are used to tighten the lower bound for the distance between two vertices. We denote the maximum of the six equations for ALP as π_t^{DL}, ALP's heuristic for A*. The following theorems then apply to π_t^{DL}:

Theorem 1: π_t^{DL} is an admissible heuristic.

Theorem 2: Using distributed embedding, π_t^{DL} is not consistent.

Theorem 3: π_t^{DL} does not dominate π_t^L over the same set of landmarks.

Theorem 4: π_t^L does not dominate π_t^{DL} over a different set of landmarks.

$d(s,t) \geq \lvert d(s,l_1) - d(l_1,l_2)\rvert - d(l_2,t)$
$d(s,t) \geq \lvert d(s,l_1) - d(l_2,t)\rvert - d(l_1,l_2)$
$d(s,t) \geq \lvert d(l_1,l_2) - d(l_2,t)\rvert - d(s,l_1)$
$d(s,t) \geq \lvert d(s,l_1) - d(l_1,t)\rvert$
$d(s,t) \geq \lvert d(s,l_2) - d(l_2,t)\rvert$
$d(s,t) \geq \frac{\lvert d(s,l_1) - d(l_1,l_2)\rvert \cdot \lvert d(l_1,l_2) - d(l_2,t)\rvert - d(s,l_1) \cdot d(l_2,t)}{d(l_1,l_2)}$

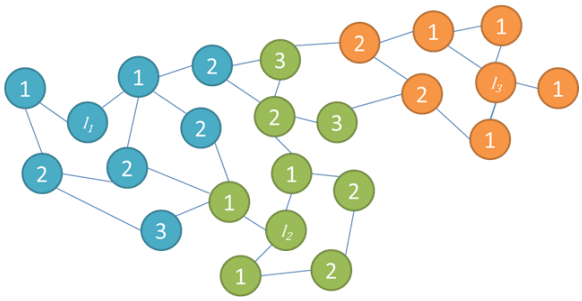

Figure 1: Dual landmark ALP heuristic inequalities(left) and a labeling example for distributed embedding on a graph (right)

Distributed Embedding

ALP's data structure can exhibit a space complexity of $\theta(|\boldsymbol{L}|^2+|V|)$ (as opposed to ALT's $\theta(|\boldsymbol{L}|\cdot|V|)$)using the following technique, called *distributed embedding*. With a partitioned graph as input, the dual landmark approach identifies a single landmark within each partition and computes a shortest path tree for the subgraph induced by each chosen landmark's graph partition. Each vertex in the graph is labeled with an identifier, signifying its landmark partition and the distance to its corresponding landmark. Any of the landmark selection methods for ALT can be used for the subgraph induced by the graph partition to select an optimal set. The final step of this process is an all-pairs shortest path calculation among the landmark set. An example of distributed embedding is illustrated in Figure 1. Allowing each landmark in the graph to access only a subset of the graph limits the size of the data structure used at query time as well as bounding the number of operations performed to actually compute the heuristic.

Results

Experimentation was performed as an initial investigation of the ALP dual landmark heuristic's behavior and its performance bounds in comparison to ALT. Comparison between ALT and ALP was performed using experimental benchmark road data from DIMACS and all available large-scale graphs (of size 10^6 nodes) that can be generated by NetworkX 1.9. Random landmark selection was used for each trial run of the two algorithms on these datasets. The Louvain algorithm [1] was used for the partitioning of each graph prior to distributed embedding. As illustrated in Figure 2, queries for paths with distances between 1 and 501 were called 10^5 times.

Discussion

We have studied the effects of using an ALP data structure of size $\theta(|\boldsymbol{L}|^2 + |V|)$ as opposed to ALT's previous $\theta(|\boldsymbol{L}| \cdot |V|)$.

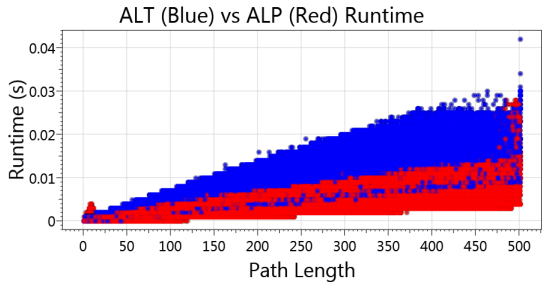

Figure 2: Graph demonstrating ALT (Blue) performance in contrast to ALP (Red) performance for growing path lengths

The resulting data show significant improvement of the runtime of the dual landmark ALP heuristic over the ALT heuristic on a diverse set of graphs with larger path lengths, as well as a significant reduction in required memory. In general, such geometric heuristics can be derived by identifying any polygon in a graph and setting the heuristic values for A* equal to the maximum derived lower bound of one of its sides. Future research will highlight the benefits and drawbacks to different landmark selection techniques for ALP and, if possible, leverage the distributed embedding method to create new forms of landmark selection that are possible only in this environment.

Acknowledgments

Special thanks to advisor Dr. Michael J. Laszlo, my doctoral committee at the Nova Southeastern University GSCIS and my employer, Raytheon BBN Technologies for their support throughout the Computer Science PhD program.

References

[1] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10, 2008.

[2] A. V. Goldberg and C. Harrelson. Computing the shortest path: A* search meets graph theory. *Proceedings of the Sixteenth Annual Acm-Siam Symposium on Discrete Algorithms*, pages 156–165, 2005.